arXiv:1608.08015v1 [cs.AI] 29 Aug 2016

Event Selection Rules to Compute Explanations

Charles Prud'homme[1], Xavier Lorca[1], and Narendra Jussien[2]

[1] TASC (CNRS/INRIA), Mines Nantes, FR – 44307 Nantes, France
FirstName.LastName@mines-nantes.fr
[2] Télécom Lille, FR – 59653 Villeneuve d'Ascq Cedex, France
Narendra.Jussien@telecom-lille.fr

Abstract. Explanations have been introduced in the previous century. Their interest in reducing the search space is no longer questioned. Yet, their efficient implementation into CSP solver is still a challenge. In this paper, we introduce ESeR, an Event Selection Rules algorithm that filters events generated during propagation. This dynamic selection enables an efficient computation of explanations for intelligent backtracking algorithms. We show the effectiveness of our approach on the instances of the last three MiniZinc challenges.

1 Introduction

In recent years, Lazy Clause Generation solvers [10] (through hybrid SAT-CSP) have brought explanations back to attention. Explanations in pure CSP solvers have suffered several drawbacks, despite their unquestioned interest: they usually required a deep refactoring of the targeted solver and computation itself was both memory and CPU consuming. In a few words, they were considered not worth the pain. In this paper, we claim that explanations have still a bright future in pure CSP solvers. We introduce here a simple yet efficient algorithm for computing explanations for CSP solvers: ESeR (Event Selection Rules). We exploit event-related information in a novel way in order to focus on those fundamentally related to a failure. Combining this technique with asynchronous constraint explanation schemas enables to easily plug intelligent backtracking algorithms in (such as CBJ [1,12]). We implemented ESeR in Choco and evaluated it on the instances of the last three MiniZinc challenges. We show its efficiency and practicability.

This paper is organized as follow. In Section 2, we recall the common background of Constraint Programming. The key elements of intelligent backtracking algorithms are introduced in Section 3. In Section 4, the pseudocode of the ESeR algorithm is presented. Comparative evaluations are showed in Section 5.

2 Background

A Constraint Satisfaction Problem (CSP) is a triple $\langle \mathcal{V}, \mathcal{D}, \mathcal{C} \rangle$ where $\mathcal{V}$ is the sequence $(v_1, v_2, \ldots, v_n)$ of *variables*, $\mathcal{D}$ is the sequence $(d_1, d_2, \ldots, d_n)$ of *domains*

associated with the variables, and $\mathcal{C}$ is the set $\{c_1, c_2, \ldots, c_m\}$ of *constraints*. The domain d_i is a finite ordered set of integer values to which the variable v_i can be assigned to. The constraint c_j, associates with a sub-sequence $\mathcal{V}(c_j) \in \mathcal{V}$, defines the allowed combinations of values which satisfy c_j. A constraint is equipped with a function p_{c_j} which removes from the domains of $\mathcal{V}(c_j)$, the values that cannot satisfy c_j. The goal of a CSP is to find a *solution*. A solution is an assignment to all the variables from their respective domains such that all the constraints are simultaneously satisfied. The exploration of the search space is processed using a *depth-first search*. At each step, a decision is taken. Without loss of generality, a (2-way) decision δ is an assignment of a variable to a value from its domain; its refutation $\neg\delta$ is the removal of this value from the variable domain. The satisfiability of all the constraints is then checked thanks to a constraint propagation algorithm which executes p_{c_j} for each constraint in turn until a fixpoint is reached (no more values can be removed). The propagation can empty a variable domain: a *failure* is detected. Upon failure, the search backtracks to the last decision and refutes it then resumes search.

How the domain of a variable changes is characterized by *events* [9,15]. An event σ is a change in a variable domain such that the cardinality of the resulting domain is strictly smaller. An event is led by a *cause*: a constraint, a decision or a refutation. There are four types of events: a removal (rem) states that a single value is removed from a variable domain, an assignment (asg) states that all values but one are removed from a variable domain, a lower bound increase (ℓow) states that all values below an (exclusive) value are removed from a variable domain, and an upper bound decrease (upp) states that all values above an (exclusive) value are removed from a variable domain. For implementation purpose, additional data is attached to each event depending on its type: x, the value removed from v, for rem, ℓ_o and ℓ_n, respectively the previous and the new lower bound of v, for ℓow, u_o and u_n, respectively the previous and the new upper bound of v, for upp, and a, ℓ_o and u_n, respectively the instantiated value of v, its previous lower and upper bound, for asg.

3 Intelligent backtracking

Nogoods and explanations have been used for a long time to improve search [3,7,8,12,14]. They enable doing intelligent backtracking (non-chronological) and can reduce significantly the number of decisions taken and speed up the resolution. Moreover, detecting which partial assignments cannot be extended to a solution avoids reproducing them in the remainder of the search tree. This is related to the notion of *nogood*.

Definition 1 (Nogood [1]). *A* nogood *is a set a of assignments that cannot all be true in any solution of the CSP.*

The definition 1 can be generalized to contain either assignments or removals.

Definition 2 (Generalized Nogood [8]). *A* g-nogood *is a set a of assignments and removals that cannot all be true in any solution of the CSP.*

Katsirelos [8] designed 1^{st}-*Decision scheme*, an algorithm to learn g-nogoods during the search. First, each event caused by a constraint is labeled with a g-nogood and the depth at which it was made, until a failure occurred. The g-nogood is the result of a direct constraint violation. Decisions are labeled with *null*. Then, exploring events in a bottom-up way iteratively modifies the g-nogood derived from the failure: the deepest event is replaced by its g-nogood, until the deepest event to replace is a decision. The resulting g-nogood is used to backtrack and to label the refutation upon backtrack. Finally, it is posted into the solver to avoid exploring the same search subtree in the future. Gent et al. [2] suggest not to store the g-nogood but a polymorphic function able to retrieve it from a given event. The g-nogood can thus be calculated lazily. They also show how such a function can be defined for various constraints. It turned out to be very effective in practice by avoiding the calculation of useless g-nogoods and thus, by saving time. Note that 1^{st}-Decision scheme does not exploit events on bounds and processes them as a sequence of removals.

A trivial nogood can be obtained from the set of decisions *explaining* the failure. This is the role of *explanations* whose definition is related to *deductions*.

Definition 3 (Deduction [4]). *A* deduction γ_σ *is the determination that an event* σ *should be applied, for instance, that a value should be removed the domain of a variable.*

Definition 4 (Explanation [4]). *An* explanation $expl(\gamma_\sigma)$ *for the deduction* γ_σ *is defined by a set of constraints* $\mathcal{C}' \subseteq \mathcal{C}$ *and a set of decisions* Δ*, such that* $\mathcal{C}' \wedge \Delta \wedge \neg\sigma$ *cannot all be true in a solution.*

Jussien and Barichard [5] introduced PaLM, an explanation based constraint programming system. In PaLM, an explanation is computed and attached when an event is generated. The code of the constraints are modified in consequence, but low-intrusive techniques can be implemented, such as the ones described in [4]. The computation may be done as and when the events are generated or by storing a list of event-related information and compiling them into explanations on demand. Then, the explanation of a failure is the union of the explanations of each value removed from the empty domain. Explanation is the key point of intelligent backtracking algorithms, such as Conflict-directed Backjumping [1,12], Dynamic Backtracking [3,6], Path Repair [7,11] or Explanation-based LNS [14]. However, the memory footprint and the CPU consumption turn out to be the major concerns of using explanations in a CSP solver. In the following, we describe an algorithm that improves the computation of explanations in a CSP solver.

4 Event selection rules to compute explanations

We now present Event SElection Rules (ESeR), an algorithm based on the ability to determine relevant event-based information. It exploits the types of

modification variables can have and relies on an asynchronous computation of explanations. We first present the pre-requisites, and then describe the notion of *event selection rule*. Finally, we describe the algorithm itself.

Pre-requisites To be able to evaluate the events, we have to store them in a chronologically ordered list, named Σ, during the propagation. Keeping the events sorted is necessary to recognize the inheritance between them. The index of the last event stored is backtrackable, thus, outdated events are automatically forgotten on backtrack. The cause which produced an event σ has an *explanation schema*, or e-schema, which is able to point out which (earlier) events may be the source of σ. A decision is event-independent and thus points out no event, a refutation is explained by previously taken decisions [4].

4.1 Event selection rules

When a domain becomes empty, the question the explanation has to answer is *what are the causes and events related to those value removals ?* In other words, how to recognize relevant events which have led, by propagation, to remove all the values from a specific domain ? To do so, we introduce *event selection rules*. An event selection rule π characterizes that a certain modification which occurred on variable v is relevant to explain the failure. It can be seen as a *negative* of an event. The type of modifications can be:

π_r**:** the event which removed a specific value from v is relevant,
π_u**:** the events which modified the lower bound of v are relevant,
π_ℓ**:** the events which modified the upper bound of v are relevant,
π_d**:** any events occurring on v are relevant.

In the event list Σ, the ℓow events (resp. *upp* events) are naturally ordered in such way that the lower of bound of a given variable only increases (resp. decreases). This is not true for rem events, that is why an integer is attached to a π_r rule representing the removed value of v, denoted $\pi_r.x$ in the following. If an event is relevant, the e-schema of the cause is interrogated. A e-schema can declare new event selection rules which will point out to earlier events. For instance, when the domain of the variable v_1 becomes empty, the first rule declared is $\langle v_1, \pi_d \rangle$ which will select all events involving v_1.

The ESeR function, described in Algorithm 1, is called when a variable domain becomes empty. The affected variable v, the list of events Σ and a boolean pe are passed as parameters of the function. By setting pe to $false$, all events from Σ are visited and a classical explanation is computed. By setting that parameter to $true$, though, the analysis of Σ stops as soon as a decision satisfies a rule and an incomplete explanation is returned. Computing classical explanations is only required for some intelligent backtracking algorithms (discussed in Section 4.2). First, an empty explanation e is created (line 2). The set of event selection rules is initialized with the rule derived from the empty domain variable (line 3). Then, a bottom-up analysis of events in Σ is run in order to find those

satisfying a rule of Π (lines 5-23). If an event satisfies a rule from Π (line 8), then the explanation of the failure needs to be updated as well as, possibly, the set of rules. These operations are executed while all the events from Σ have not been visited or a stop condition is encountered (line 6). When the iteration ends and if incomplete explanations are allowed, the rules gathered up to this point are copied into the explanation (line 24-25). In this way, the computation may be resumed, if necessary (see Section 4.2). Note that, merging two explanations (line 13) also merges rules, if any. The computed explanation is then returned (line 27).

Algorithm 1 ESeR: Event Selection Rule

```
 1: function EXPLAIN(an empty domain variable: v, a list of events: Σ, boolean pe)
 2:     e ← ∅                                          ▷ Create a new explanation
 3:     Π ← ⟨v, π_d⟩                                   ▷ Initialize the rules set
 4:     stop ← false
 5:     i = |Σ| − 1
 6:     while i ≥ 0 and ¬stop do
 7:         σ ← Σ[i]
 8:         if σ satisfies a rule of Π then            ▷ See Table 1
 9:             if σ.c is a decision then
10:                 add the decision to e
11:                 stop ← pe
12:             else if σ.c is a refutation then
13:                 merge the explanation of the refutation into e
14:             else                                   ▷ The cause is a constraint
15:                 add the constraint to e
16:                 update Π with rules returned by the constraint's e-schema
17:             end if
18:             if σ.t = rem then
19:                 remove the corresponding rule from Π
20:             end if
21:         end if
22:         i = i − 1
23:     end while
24:     if pe then
25:         copy rules in e
26:     end if
27:     return e
28: end function
```

The satisfaction test (Algorithm 1, line 8) simply verifies if a rule based on the event's variable exists and if the event's modification type is covered by the rule's type, which may imply an evaluation of $\sigma.X$ and $\pi_r.x$. The covering conditions are depicted in Table 1. For instance, a π_d selection rule exists for a variable v_1 covers a ℓow event for v_1.

Table 1. Covering conditions to select an event.

		event's type			
		asg	ℓow	upp	rem
rule's type	π_d	$true$	$true$	$true$	$true$
	$\pi_\ell \wedge \pi_u$	$true$	$true$	$true$	$\sigma.r \notin d_{\sigma.v}$
	π_ℓ	$\sigma.r \notin d_{\sigma.v}$	$true$	$false$	$\sigma.r \notin d_{\sigma.v}$
	π_u	$\sigma.r \notin d_{\sigma.v}$	$false$	$true$	$\sigma.r \notin d_{\sigma.v}$
	π_r	$\pi_r.x \in [\![\sigma.\ell_o, \sigma.u_o]\!]$	$\pi_r.x \in [\![\sigma.u_n, \sigma.u_o]\!]$	$\pi_r.x \in [\![\sigma.\ell_o, \sigma.\ell_n]\!]$	$\pi_r.x = \sigma.r$

When updating the explanation (Algorithm 1, line 9-21), three cases are considered: (a) the cause is a decision (line 9-11), (b) the cause is a refutation (line 12-13) and (c) the cause is a constraint (line 15-16). In (a), the decision is added to e. Then, if an incomplete explanation is allowed, *stop* is set to *true* in order to interrupt the iteration over events. In (b), a previous failure led to refute a decision whom explanation was previously computed. That explanation is then merged into the current one. In (c) the constraint is added to e. Then, the e-schema of the constraint is interrogated to add new rules to Π. A valid e-schema for all constraint is to declare a π_d rule for all variables of the constraint. However, it results in weak explanation, and specific e-schemas should be implemented for each constraint. Finally, the set of rules Π can be reduced whenever a π_r rule is satisfied, since it can appear only once in a search branch (line 18-20).

The ESeR algorithm runs in $\mathcal{O}(|\Sigma| \times s)$ where s is complexity of the most complex e-schema, at each failure. Indeed, the while-loop (line 8-23) only executes $\mathcal{O}(1)$ operations except when it interrogates the explanation schema of the cause (line 16). The complexity of determining which rules need to be added to the rules set depends on the type of cause.

4.2 Using ESeR in intelligent backtracking algorithms

Intelligent backtracking algorithms can be split in three groups. First, those requiring classical explanations, such as Path Repair or Explanation-based LNS. They rely on a relaxation of the explaining decisions set, so all of them must be retrieved to be correct. Second, those accepting incomplete explanations such as Conflict-directed Backjumping. Only the deepest decision from the explaining ones is interesting to jump back to it, so finding it is enough. Finally, those accepting incomplete explanations with the necessity to keep up their computation later, such as Dynamic Backtracking. Dynamic Backtracking mimics Conflict-directed Backjumping, but when jumping the deepest decision, decisions and refutations taken between the failure and the one to jump back to have to be kept by the search heuristic. Refutations depend on earlier decisions and, potentially, on the one to jump back to. Stating whether or not a refutation depends on the jump back decision is made by keeping up their computation until the

event corresponding to that decision is visited. If the decision satisfies one of the rules, the refutation depends on the jump back decision.

Implementation For implementation concerns, all rules involving the same variable are represented with a unique object made of two integers and an integers set. The two integers correspond to the variable index and the aggregated modification types (handled with bitwise operations). The integers set manages removed values. An explanation is composed of a decisions set (encoded in a bitset), a constraints set and a rules set, possibly empty. Embedding a rules set in the explanation is motivated to keep up computing it, which is required when dealing with incomplete explanations. Finally, the rules maintained in the ESeR algorithm can be compiled in order to extract g-nogoods. For efficiency purpose, the g-nogood must be reviewed to deal with bound modifications, in such way that they are not handled as a sequence of value removals but as native cases. Extracting g-nogoods from the rules set is not implemented in the evaluated version of ESeR.

5 Evaluation

Our explanation engine has been implemented in Choco-3.3.1[3] [13], a Java library for constraint programming. We evaluated three configurations: a standard backtrack search (`choco`), a conflict-directed backjumping search disabling incomplete explanations (`CBJ`) and CBJ enabling incomplete explanations (`CBJ-i`). Not all (global) constraints are natively explained in Choco. If not, the naive but valid e-schema depicted in Section 4.1 is used. The set of benchmarks used consists of the instances from the MiniZinc Challenges[4] (2012, 2013 and 2014), and it is composed of 269 instances from 66 problems. The models rely on a large variety of constraints, including global constraints, and each problem describes a dedicated search strategy. Each of the 269 instances was executed with a 15-minute timeout, on a Macbook Pro with 6-core Intel Xeon at 2.93Ghz running a MacOS 10.10, and Java 1.8.0_05. Each instance was run on its own core, each with up to 4096MB of memory. Note that because there are too many distinct types of problems (66), we did not characterize them on the scatterplots.

Figure 1 reports the comparative evaluations: (a) `CBJ` vs. `choco` on the left plot and (b) `CBJ` vs. `CBJ-i` on the right plot. Note that the axes are in logarithmic scale. On plot (a), the contribution of using explanations on some certain instances is clear, it corresponds to the points above the diagonal. Even if almost 50% of instances timed out for both approaches, `CBJ` is faster in 46% of the remaining instances, and the speedups can go up to 286x. On plot (b), the benefit of using incomplete explanations is without any doubts in favor of `CBJ-i`: most of the points are below the diagonal. Here again, up to 50% of the instances timed out for the two approaches but enabling incomplete explanation is particularly remarkable: 79% of the remaining instances are solved faster with `CBJ-i`.

[3] Available on `www.choco-solver.org`

[4] `http://www.minizinc.org/challenge.html`

The speedups can go up to 47x. We do not report the comparative evaluations of `choco` and Dynamic Backtracking, but the results (available on request) are very close to those observed with CBJ.

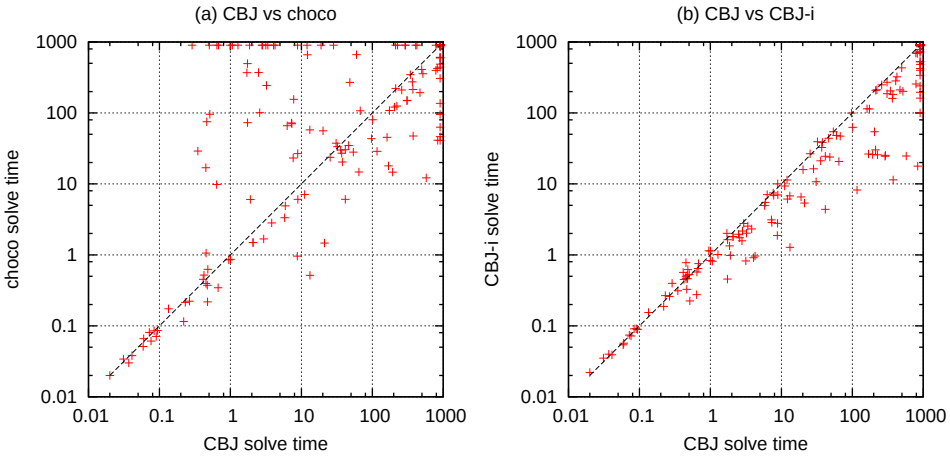

Fig. 1. Comparative runtime evaluations. Each point corresponds to a single instance. The x-axis is the solve time (in seconds) for a first approach, the y-axis is the solve time (in seconds) for the second one. A point above the diagonal means the first approach was faster, below the diagonal that the second one was faster.

6 Conclusion

In this paper, we introduced ESeR, an Event Selection Rules algorithm for CSP solvers that filters events generated during propagation. This dynamic selection enables an efficient computation of explanations for intelligent backtracking algorithms. The efficiency of our approach has been validated, in practice, on a dataset of 269 instances from the three last MiniZinc challenges. The further works include an improvement of the possibilities offered by ESeR, by extracting g-nogoods from the selection rules, and a study of alternative intelligent backtracking algorithms.

References

1. Rina Dechter. Enhancement schemes for constraint processing: Backjumping, learning, and cutset decomposition. *Artif. Intell.*, 41(3):273–312, 1990.
2. Ian P. Gent, Ian Miguel, and Neil C. A. Moore. Lazy explanations for constraint propagators. In Manuel Carro and Ricardo Peña, editors, *Practical Aspects of Declarative Languages, 12th International Symposium, PADL 2010, Madrid, Spain, January 18-19, 2010. Proceedings*, volume 5937 of *Lecture Notes in Computer Science*, pages 217–233. Springer, 2010.

3. Matthew L. Ginsberg. Dynamic backtracking. *J. Artif. Intell. Res. (JAIR)*, 1:25–46, 1993.
4. Narendra Jussien. *The versatility of using explanations within constraint programming*. Habilitation à diriger des recherches, Université de Nantes, September 2003.
5. Narendra Jussien and Vincent Barichard. The palm system: explanation-based constraint programming. In *In Proceedings of TRICS: Techniques foR Implementing Constraint programming Systems, a post-conference workshop of CP 2000*, pages 118–133, 2000.
6. Narendra Jussien, Romuald Debruyne, and Patrice Boizumault. Maintaining arc-consistency within dynamic backtracking. In *Principles and Practice of Constraint Programming (CP 2000)*, number 1894 in Lecture Notes in Computer Science, pages 249–261, Singapore, September 2000. Springer-Verlag.
7. Narendra Jussien and Olivier Lhomme. The path repair algorithm. *Electronic Notes in Discrete Mathematics*, 4:2–16, 2000.
8. George Katsirelos. *Nogood Processing in CSPs*. PhD thesis, University of Toronto, 2009.
9. F. Laburthe. Choco: implementing a cp kernel. In *Proceedings of Techniques foR Implementing Constraint programming Systems (TRICS'00)*, pages 118–133, 2000.
10. Olga Ohrimenko, Peter J. Stuckey, and Michael Codish. Propagation via lazy clause generation. *Constraints*, 14(3):357–391, 2009.
11. Cédric Pralet and Gérard Verfaillie. Travelling in the world of local searches in the space of partial assignments. In *CPAIOR*, pages 240–255, 2004.
12. P. Prosser. MAC-CBJ: maintaining arc consistency with conflict-directed backjumping. Technical Report Research Report/95/177, Dept. of Computer Science, University of Strathclyde, 1995.
13. Charles Prud'homme, Jean-Guillaume Fages, and Xavier Lorca. *Choco3 Documentation*. TASC, INRIA Rennes, LINA CNRS UMR 6241, COSLING S.A.S., 2014.
14. Charles Prud'homme, Xavier Lorca, and Narendra Jussien. Explanation-based large neighborhood search. *Constraints*, 19(4):339–379, 2014.
15. Christian Schulte and Peter J. Stuckey. Efficient constraint propagation engines. *ACM Trans. Program. Lang. Syst.*, 31(1), 2008.